

Maximizing Expected Utility for Stochastic Combinatorial Optimization Problems

Jian Li^{*1} and Amol Deshpande^{†2}

¹Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, P.R.China

²Department of Computer Science, University of Maryland, College Park, USA

Abstract

We study the stochastic versions of a broad class of combinatorial problems where the weights of the elements in the input dataset are uncertain. The class of problems that we study includes shortest paths, minimum weight spanning trees, and minimum weight matchings over probabilistic graphs, and other combinatorial problems like knapsack. We observe that the expected value is inadequate in capturing different types of *risk-averse* or *risk-prone* behaviors, and instead we consider a more general objective which is to maximize the *expected utility* of the solution for some given utility function, rather than the expected weight (expected weight becomes a special case). We show that we can obtain a polynomial time approximation algorithm with *additive error* ϵ for any $\epsilon > 0$, if there is a pseudopolynomial time algorithm for the *exact* version of the problem (This is true for the problems mentioned above) and the maximum value of the utility function is bounded by a constant. ¹ Our result generalizes several prior results on stochastic shortest path, stochastic spanning tree, and stochastic knapsack. Our algorithm for utility maximization makes use of the separability of exponential utility and a technique to decompose a general utility function into exponential utility functions, which may be useful in other stochastic optimization problems.

1 Introduction

The most common approach to deal with optimization problems in presence of uncertainty is to optimize the expected value of the solution. However, expected value is inadequate in expressing diverse people's preferences towards decision-making under uncertain scenarios. In particular, it fails at capturing different *risk-averse* or *risk-prone* behaviors that are commonly observed. Consider the following simple example where we have two lotteries L_1 and L_2 . In L_1 , the player could win 1000 dollars with probability 1.0, while in L_2 the player could win 2000 dollars with probability 0.5 and 0 dollars otherwise. It is easy to see that both have the same expected payoff of 1000 dollars. However, many, if not most, people would treat L_1 and L_2 as two completely different choices. Specifically, a risk-averse player is likely to choose L_1 and a risk-prone player may prefer L_2 (Consider a gambler who would like to spend 1000 dollars to play double-or-nothing). A more involved but also more surprising example is the *St. Petersburg paradox* (see

^{*}lapordge@gmail.com

[†]amol@cs.umd.edu

¹Following the literature [46], we differentiate between *exact* version and *deterministic* version of a problem; in the exact version of the problem, we are given a target value and asked to find a solution (e.g., a path) with exactly that value (i.e., path length).

e.g., [37, 1]) which has been widely used in the economics literature as a criticism of expected value. See Appendix A for a brief description of the St. Petersburg paradox. These observations and criticisms have led researchers, especially in Economics, to study the problem from a more fundamental perspective and to directly maximize user satisfaction, often called *utility*. The uncertainty present in the problem instance naturally leads us to optimize the *expected utility*.

Let \mathcal{F} be the set of feasible solutions to an optimization problem. Each solution $S \in \mathcal{F}$ is associated with a random weight $w(S)$. For instance, \mathcal{F} could be a set of lotteries and $w(S)$ is the (random) payoff of lottery S . We model the risk awareness of a user by a utility function $\mu : \mathbb{R} \rightarrow \mathbb{R}$: the user obtains $\mu(x)$ units of utility if the outcome is x , i.e., $w(S) = x$. Formally, the *expected utility maximization principle* is simply stated as follows: the most desirable solution S is the one that maximizes the expected utility, i.e.,

$$S = \arg \max_{S' \in \mathcal{F}} \mathbb{E}[\mu(w(S'))]$$

Indeed, the expected utility theory is a branch of the utility theory that studies “betting preferences” of people with regard to uncertain outcomes (gamblers). The theory was formally initiated by von Neumann and Morgenstern in 1940s [55, 21]² who gave an axiomatization of the theory (known as *von Neumann-Morgenstern expected utility theorem*). The theory is well known to be versatile in expressing diverse risk-averse or risk-prone behaviors.

In this paper, we focus on the following broad class of combinatorial optimization problems. The deterministic version of the problem has the following form: we are given a ground set of elements $U = \{e_i\}_{i=1\dots n}$; each element e is associated with a weight w_e ; each feasible solution is a subset of the elements satisfying some property. Let \mathcal{F} denote the set of feasible solutions. The objective for the deterministic problem is to find a feasible solution S with the minimum total weight $w(S) = \sum_{e \in S} w_e$. We can see that many combinatorial problems such as shortest path, minimum spanning tree, and minimum weight matching belong to this class. In the stochastic version of the problem, the weight w_e of each element e is a nonnegative random variable. We assume all w_e s are independent of each other. We use $p_e(\cdot)$ to denote the probability density function for w_e (or probability mass function in discrete case). We are also given a utility function $\mu : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ which maps a weight value to a utility value. By the expected utility maximization principle, our goal here is to find a feasible solution $S \in \mathcal{F}$ that maximizes the expected utility, i.e., $\mathbb{E}[\mu(w(S))]$. We call this problem the *expected utility maximization (EUM)* problem.

Let us use the following toy example to illustrate the rationale behind EUM. There is a graph with two nodes s and t and two parallel links e_1 and e_2 . Edge e_1 has a fixed length 1 while the length of e_2 is 0.9 with probability 0.9 and 1.9 with probability 0.1 (the expected value is also 1). We want to choose one edge to connect s and t . It is not hard to imagine that a risk-averse user would choose e_1 since e_2 may turn out to be a much larger value with a nontrivial probability. We can capture such behavior using the utility function (1) (defined in Section 1.1). Similarly, we can capture the risk-prone behavior by using, for example, the utility function $\mu(x) = \frac{1}{x+1}$. It is easy to see that e_1 maximizes the expected utility in the former case, and e_2 in the latter.

1.1 Our Contributions

We discuss in detail our result for EUM. We assume μ is part of the specification of the problem but not part of the input. Moreover, we assume $\lim_{x \rightarrow \infty} \mu(x) = 0$. This captures the fact that if the weight of solution is too large, it becomes almost useless for us. W.l.o.g. we can also assume $0 \leq \mu(x) \leq 1$ for $x \geq 0$, by scaling.

²Daniel Bernoulli also developed many ideas, such as risk aversion and utility, in his work *Specimen theoriae novae de mensura sortis* (*Exposition of a New Theory on the Measurement of Risk*) in 1738 [8].




Figure 1: (1) The utility function $\tilde{\chi}(x)$, a continuous variant of the threshold function $\chi(x)$; (2) A smoother variant of $\chi(x)$; (3) The utility function $\tilde{\chi}_2(x)$, a continuous variant of the 2-d threshold function $\chi_2(x)$.

We say a function $\tilde{\mu}(x)$ is an ϵ -approximation of $\mu(x)$ if $|\tilde{\mu}(x) - \mu(x)| \leq \epsilon \forall x \geq 0$. For ease of exposition, we let $\tilde{\mu}(x)$ be a complex function. Recall that a polynomial time approximation scheme (PTAS) is an algorithm which takes an instance of a minimization problem and a parameter ϵ and produces a solution whose cost is within a factor $1 + \epsilon$ of the optimum, and the running time, for any fixed ϵ , is polynomial in the size of the input. We use A to denote the deterministic combinatorial optimization problem under consideration. The *exact version* of a problem A asks the question whether there is a feasible solution of A with weight exactly equal to a given number K . We say an algorithm runs in *pseudopolynomial time* for the exact version of A if the running time is polynomial in n and K . Our first main theorem is the following.

Theorem 1 Assume that there is a pseudopolynomial algorithm for the exact version of A . Further assume that given any $\epsilon > 0$, we can find an ϵ -approximation of the utility function μ as $\tilde{\mu}(x) = \sum_{k=1}^L c_k \phi_k^x$, where L is a constant and $|\phi_k| \leq 1 \forall k$; ϕ_k may be complex numbers. Then, there is an algorithm that runs in time $(n/\epsilon)^{O(L)}$ that approximates $\text{EUM}(A)$ with an additive error $O(\epsilon)$. If the optimal expected utility is $\Theta(1)$, we obtain a PTAS.

For many combinatorial problems, a pseudopolynomial algorithm for the exact version is known. Examples include shortest path, spanning tree, matching and knapsack. Hence, the only task left is to find a short exponential sum that ϵ -approximates μ . For this purpose, we adopt the Fourier series technique. However, the technique cannot be used directly since it works only for periodic functions with bounded periodicities. In order to get a good approximation for $x \in [0, \infty)$, we leverage the fact that $\lim_{x \rightarrow \infty} \mu(x) = 0$ and develop a general framework that uses the Fourier series decomposition as a subroutine. Generally speaking, such an approximation is only possible if the function is “well behaved”, i.e., it satisfies some continuity or smoothness conditions. In particular, we prove Theorem 2. We say that the utility function μ satisfies the α -Hölder condition if $|\mu(x) - \mu(y)| \leq C|x - y|^\alpha$, for some constant C and some constant α .

Theorem 2 If μ satisfies the α -Hölder condition for some constant $\alpha > 1/2$, then, for any $\epsilon > 0$, we can obtain an exponential sum with $O(\text{poly}(\frac{1}{\epsilon}))$ terms which is an ϵ -approximation of μ for $x \geq 0$.

Consider the utility function

$$\tilde{\chi}(x) = \begin{cases} 1 & x \in [0, 1] \\ -\frac{x}{\delta} + \frac{1}{\delta} + 1 & x \in [1, 1 + \delta] \\ 0 & x > 1 + \delta \end{cases} \quad (1)$$

where $\delta > 0$ is a small constant (See Figure 1(1)). We can verify that $\tilde{\chi}$ satisfies 1-Hölder condition with $C = \frac{1}{\delta}$. Therefore, Theorem 2 is applicable. This example is interesting since it can be viewed as a continuous variant of the threshold function

$$\chi(x) = \begin{cases} 1 & x \in [0, 1] \\ 0 & x > 1 \end{cases}, \quad (2)$$

for which maximizing the expected utility is equivalent to maximizing $\Pr(w(S) \leq 1)$. This special case has been considered several times in literature for various problems including stochastic shortest path [43], stochastic spanning tree [30, 23], stochastic knapsack [24] and some other stochastic problems [3, 41].

It is interesting to compare our result with the result for the stochastic shortest path problem considered by Nikolova et al. [43, 41]. In [43], they show that there is an exact $O(n^{\log n})$ time algorithm for maximizing the probability that the length of the path is at most 1, i.e., $\Pr(w(S) \leq 1)$, assuming all edges are normally distributed and there is a path with its mean at most 1. Later, Nikolova [41] extends the result to an FPTAS for any problem under the same assumptions, if the deterministic version of the problem has a polynomial time exact algorithm. We can see that under such assumptions, the optimal probability is at least $1/2$.³ Therefore, provided the same assumption and further assuming that $\Pr(w_e < 0)$ is miniscule,⁴ our algorithm is a PTAS for the continuous variant of the problem. Indeed, we can translate this result to a bi-criterion approximation result of the following form: for any fixed $\delta, \epsilon > 0$, we can find in polynomial time a solution S such that

$$\Pr(w(S) \leq 1 + \delta) \geq (1 - \epsilon) \Pr(w(S^*) \leq 1).$$

where S^* is the optimal solution (Corollary 2). We note that such a bi-criterion approximation was only known for exponentially distributed edges before [43].

Let us consider another application of our results to the stochastic knapsack problems defined in [24]. Given a set U of independent random variables $\{x_1, \dots, x_n\}$, with associated profits $\{v_1, \dots, v_n\}$ and an overflow probability γ , we are asked to pick a subset S of U such that

$$\Pr\left(\sum_{i \in S} x_i \geq 1\right) \leq \gamma$$

and the total profit $\sum_{i \in S} v_i$ is maximized. Goel and Indyk [24] showed that, for any $\epsilon > 0$, there is a polynomial time algorithm that can find a solution S with the profit as least the optimum and $\Pr(\sum_{i \in S} x_i \geq 1 + \epsilon) \leq \gamma(1 + \epsilon)$ for exponentially distributed variables. They also gave a quasi-polynomial time approximation scheme for Bernoulli distributed random variables. Quite recently, in parallel with our work, Bhalgat et al. [13] obtained the same result for arbitrary distributions under the assumption that $\gamma = \Theta(1)$. Their technique is based on discretizing the distributions and is quite involved. Our result, applied to stochastic knapsack, matches that of Bhalgat et al. We remark that our algorithm is much simpler and has a much better running time (Theorem 5). Despite a little loss in the approximation guarantees in some cases, our technique can be applied to almost all positive probability distributions, and a much richer class of utility functions.

Equally importantly, we can extend our basic approximation scheme to handle generalizations such as multiple utility functions and multidimensional weights. Interesting applications of these extensions include generalizations of stochastic knapsack, such as *stochastic multiple knapsack* (Theorem 8) and *stochastic multidimensional knapsack (stochastic packing)* (Theorem 9).

1.2 Related Work

In recent years stochastic optimization problems have drawn much attention from the computer science community and stochastic versions of many classical combinatorial optimization problems have been studied. In particular, a significant portion of the efforts has been devoted to the two-stage stochastic optimization

³The sum of multiple Gaussians is also a Gaussian. Hence, if we assume the mean of the length of a path (which is a Gaussian) is at most 1, the probability that the length of the path is at most 1 is at least $1/2$.

⁴Our technique can only handle distributions with positive supports. Thus, we have to assume that the probability that a negative value appears is miniscule and can be safely ignored.

problem. In such a problem, in a first stage, we are given probabilistic information about the input but the cost of selecting an item is low; in a second stage, the actual input is revealed but the costs for the elements are higher. We are asked to make decision after each stage and minimize the expected cost. Some general techniques have been developed [28, 50]. We refer interested reader to [54] for a comprehensive survey. Another widely studied type of problems considers designing *adaptive* probing policies for stochastic optimization problems where the existence or the exact weight of an element can be only known upon a probe. There is typically a budget for the number of probes (see e.g., [27, 18]), or we require an irrevocable decision whether to include the probed element in the solution right after the probe (see e.g., [20, 16, 5, 19, 13]). However, most of those works focus on optimizing the expected value of the solution. There is also sporadic work on optimizing the overflow probability or some other objectives subject to the overflow probability constraints. In particular, a few recent works have explicitly motivated such objectives as a way to capture the risk-averse type of behaviors [3, 41, 53]. Besides those works, there has been little work on optimizing more general utility functions for combinatorial stochastic optimization problems from an approximation algorithms perspective.

The most related work to ours is the stochastic shortest path problem (Stoch-SP), which was also the initial motivation for this work. The problem has been studied extensively for several special utility functions in operation research community. Sigal et al. [51] studied the problem of finding the path with greatest probability of being the shortest path. Loui [36] showed that Stoch-SP reduces to the shortest path (and sometimes longest path) problem if the utility function is linear or exponential. Nikolova et al. [42] identified more specific utility and distribution combinations that can be solved optimally in polynomial time. Much work considered dealing with more general utility functions, such as piecewise linear or concave functions, e.g., [39, 40, 7]. However, these algorithms are essentially heuristics and the worst case running times are still exponential. Nikolova et al. [43] studied the problem of maximizing the probability that the length of the chosen path is less than some given parameter. Besides the result we mentioned before, they also considered Poisson and exponential distributions. Despite much effort on this problem, no algorithm is known to run in polynomial time and have provable performance guarantees, especially for more general utility functions or more general distributions. This is perhaps because the hardness comes from different sources, as also noted in [43]: the shortest path selection per se is combinatorial; the distribution of the length of a path is the convolution of the distributions of its edges; the objective is nonlinear; to list a few.

Kleinberg et al. [32] first considered the stochastic knapsack problem with Bernoulli-type distributions and provided a polynomial-time $O(\log 1/\gamma)$ approximation where γ is the given overflow probability. For item sizes with exponential distributions, Goel and Indyk [24] provided a bi-criterion PTAS, and for Bernoulli-distributed items they gave a quasi-polynomial approximation scheme. Chekuri and Khanna [15] pointed out that a PTAS can be obtained for the Bernoulli case using their techniques for the multiple knapsack problem. Goyal and Ravi [26] showed a PTAS for Gaussian distributed sizes. Quite recently, Bhalgat, Goel and Khanna [13] developed a general discretization technique that reduces the distributions to a small number of equivalent classes which we can efficiently enumerate for both adaptive and nonadaptive versions of stochastic knapsack. They used this technique to obtain improved results for several variants of stochastic knapsack, notably a bi-criterion PTAS for the *adaptive version* of the problem. Dean et al. [20] gave the first constant approximation for the adaptive version of stochastic knapsack. The adaptive version of stochastic multidimensional knapsack (or equivalently stochastic packing) has been considered in [19, 13] where constant approximations and a bi-criterion PTAS were developed.

This work is partially inspired by our prior work on top- k and other queries over probabilistic datasets [33, 35]. In fact, we can show that both the *consensus answers* proposed in [33] and the *parameterized ranking functions* proposed in [35] follow the expected utility maximization principle where the utility functions

are materialized as distance metrics for the former and the weight functions for the latter. Our technique for approximating the utility functions is also similar to the approximation scheme used in [35] in spirit. However, no performance guarantees are provided in that work.

There is a large volume of work on approximating functions using short exponential sums over a bounded domain, e.g., [45, 9, 10, 11]. Some works also consider using linear combinations of Gaussians or other *kernels* to approximate functions with finite support over the entire real axis $(-\infty, +\infty)$ [17]. This is however impossible using exponentials since α^x is either periodic (if $|\alpha| = 1$) or approaches to infinity when $x \rightarrow +\infty$ or $x \rightarrow -\infty$ (if $|\alpha| \neq 1$).

2 Algorithm

We first note that EUM is #P-hard in general since the problem of computing the overflow probability of a set of items with Bernoulli distributions, a very special case of our problem, is #P-hard [32].

Our approach is very simple. We first observe that the problem is easy if the utility function is an exponential function. We approximate the utility function $\mu(x)$ by a short exponential sum, i.e., $\sum_{i=1}^L c_i \phi_i^x$ with L being a constant (c_i and ϕ_i may be complex numbers). Hence, $\mathbb{E}[\mu(w(S))]$ can be approximated by $\sum_{i=1}^L c_i \mathbb{E}[\phi_i^{w(S)}]$. Then, we consider the following multi-criterion version of the problem with L objectives $\{\mathbb{E}[\phi_i^{w(S)}]\}_{i=1,\dots,L}$: given L complex numbers v_1, \dots, v_L , we want to find a solution S such that $\mathbb{E}[\phi_i^{w(S)}] \approx v_i$ for $i = 1, \dots, L$. We achieve this by utilizing the pseudopolynomial time algorithm for the exact version of the problem. We argue that we only need to consider a polynomial number of v_1, \dots, v_L combinations (which we call *configurations*) to find out the approximate optimum. In Section 2.1, we show how to solve the multi-criterion problem provided that a short exponential sum approximation of μ is given. In particular, we prove Theorem 1. Then, we show how to approximate μ by a short exponential sum by proving Theorem 2 in Section 2.2 and Section 2.3.

Let us first consider the exponential utility function $\mu(x) = \alpha^x$ for any $\alpha \in \mathbb{C}$. Fix an arbitrary solution S and $\alpha > 0$. Due to the independence of the elements, we can see that

$$\mathbb{E}[\alpha^{w(S)}] = \mathbb{E}[\alpha^{\sum_{e \in S} w_e}] = \mathbb{E}[\prod_{e \in S} \alpha^{w_e}] = \prod_{e \in S} \mathbb{E}[\alpha^{w_e}]$$

Taking log on both sides, we get $\log \mathbb{E}[\alpha^{w(S)}] = \sum_{e \in S} \log \mathbb{E}[\alpha^{w_e}]$. If α is a positive real number and $\mathbb{E}[\alpha^{w_e}] \leq 1$ (or equivalently, $-\log \mathbb{E}[\alpha^{w_e}] \geq 0$), this reduces to the deterministic optimization problem.

We still need to show how to compute $\mathbb{E}[\alpha^{w_e}]$. If w_e is a discrete random variable with a polynomial size support, we can easily compute $\mathbb{E}[\alpha^{w_e}]$ in polynomial time. If w_e has an infinite discrete or continuous support, we can not compute $\mathbb{E}[\alpha^{w_e}]$ directly and may need to approximate it. We briefly discuss this issue and its implications on our results in Appendix C.

2.1 Proof of Theorem 1

Now, we prove Theorem 1. We start with some notations. We use $|c|$ and $\arg(c)$ to denote the absolute value and the argument of the complex number c , respectively. In other words, $c = |c|(\cos(\arg(c)) + i \sin(\arg(c))) = |c|e^{i \arg(c)}$. We always require $\arg(c) \in [0, 2\pi)$ for any $c \in \mathbb{C}$. Recall that we say the exponential sum $\sum_{i=1}^L c_i \phi_i^x$ is an ϵ -approximation for $\mu(x)$ if the following holds:

$$|\mu(x) - \sum_{i=1}^L c_i \phi_i^x| \leq \epsilon \quad \forall x \geq 0$$

We first show that if the utility function can be decomposed exactly into a short exponential sum, we can approximate the optimal expected utility well.

Theorem 3 Assume $\tilde{\mu}(x) = \sum_{k=1}^L c_k \phi_k^x$ is the utility function where $|\phi_k| \leq 1$ for $1 \leq k \leq L$. We also assume that there is a pseudopolynomial algorithm for the exact version of A . Then, for any $\epsilon > 0$, there is an algorithm that runs in time $(n/\epsilon)^{O(L)}$ and finds a solution S such that

$$|\mathbb{E}[\tilde{\mu}(w(S))] - \mathbb{E}[\tilde{\mu}(w(\tilde{S}))]| < \epsilon$$

where $\tilde{S} = \arg \max_{S'} |\mathbb{E}[\tilde{\mu}(w(S'))]|$.

We use the scaling and rounding technique that has been used often in multi-criterion optimization problems (e.g., [49, 46]). Since our objective function is not additive and not monotone, the general results for multi-criterion optimization [46, 38, 49, 2] do not directly apply here. We briefly sketch our algorithm. Let $\gamma = \delta = \frac{\epsilon}{Ln}$. For each $e \in U$, we associate it with a $2L$ dimensional integer vector

$$\langle a_1(e), b_1(e), \dots, a_L(e), b_L(e) \rangle \text{ where } a_i(e) = \lfloor \frac{-\ln |\mathbb{E}[\phi_i^{w_e}]|}{\gamma} \rfloor \text{ and } b_i(e) = \lfloor \frac{\arg(\mathbb{E}[\phi_i^{w_e}])}{\delta} \rfloor.$$

$a_i(e)$ and $b_i(e)$ are the scaled and rounded versions of $-\ln |\mathbb{E}[\phi_i^{w_e}]|$ and $\arg(\mathbb{E}[\phi_i^{w_e}])$, respectively. Since $|\phi_i| \leq 1$, we can see that $a_i(e) \geq 0$ for any $e \in U$. We maintain $(JK)^L$ configurations where $J = \lceil \frac{-\ln(\epsilon/L)}{\gamma} \rceil$ and $K = \lceil \frac{2\pi n}{\delta} \rceil$. The number of configurations is $(n/\epsilon)^{O(L)}$. Each configuration $\sigma(\mathbf{a})$ is indexed by a $2L$ -dimensional vector $\mathbf{a} = \langle \alpha_1, \beta_1, \dots, \alpha_L, \beta_L \rangle$ where $1 \leq \alpha_i \leq J$ and $1 \leq \beta_i \leq K$ for $i = 1, \dots, L$. In other words, the configurations are $\sigma(\langle 1, 1, \dots, 1, 1 \rangle), \dots, \sigma(\langle J, K, \dots, J, K \rangle)$. For vector $\mathbf{a} = \langle \alpha_1, \beta_1, \dots, \alpha_L, \beta_L \rangle$, configuration $\sigma(\mathbf{a}) = 1$ if and only if there is a feasible solution $S \in \mathcal{F}$ such that for all $j = 1, \dots, L$, $\beta_j = \sum_{e \in S} b_j(e)$, and $\alpha_j = \min(J, \sum_{e \in S} a_j(e))$. Otherwise, $\sigma(\mathbf{a}) = 0$. Lemma 1 tells us the expected utility for the rounded instance is close to the true value of the expected utility. Lemma 2 shows we can compute those configurations in polynomial time.

Lemma 1 For vector $\mathbf{a} = \langle \alpha_1, \beta_1, \dots, \alpha_L, \beta_L \rangle$, $\sigma_v(\mathbf{a}) = 1$ if and only if there is a solution S such that

$$|\mathbb{E}[\tilde{\mu}(w(S))] - \sum_{k=1}^L c_k e^{-\alpha_k \gamma + i \beta_k \delta}| \leq O(\epsilon).$$

Lemma 2 Suppose there is a pseudopolynomial time algorithm for the exact version of A , which runs in time polynomial in n and t (t is the maximum integer in the instance of A). Then, we can compute the values for these configurations in time $(\frac{n}{\epsilon})^{O(L)}$.

The missing proofs can be found in Appendix B. Now, we can easily prove Theorem 3.

Proof of Theorem 3: We first use the algorithm in Lemma 2 to compute the values for all configurations. Then, we find the configuration $\sigma(\langle \alpha_1, \beta_1, \dots, \alpha_L, \beta_L \rangle)$ that has value 1 and that maximizes the quantity $|\sum_{k=1}^L c_k e^{-\alpha_k \gamma + i \beta_k \delta}|$. The feasible solution S corresponding to this configuration is our final solution. It is easy to see that the theorem follows from Lemma 1. \square

Theorem 1 can be readily obtained from Theorem 3 and the fact $\tilde{\mu}$ is an ϵ -approximation of μ .

Proof of Theorem 1: Suppose S is our solution and S^* is the optimal solution for utility function μ . From Theorem 3, we know that $|\mathbb{E}[\tilde{\mu}(w(S))] - \mathbb{E}[\tilde{\mu}(w(S^*))]| \leq \epsilon$. Since $\tilde{\mu}$ is an ϵ -approximation of μ , we can see that

$$|\mathbb{E}[\mu(w(S))] - \mathbb{E}[\tilde{\mu}(w(S))]| = \left| \int (\mu(x) - \tilde{\mu}(x))p_S(x)dx \right| \leq \left| \int \epsilon p_S(x)dx \right| \leq \epsilon$$

for any solution S , where p_S is the probability density function of S . Therefore, we have

$$|\mathbb{E}[\mu(w(S))]| \geq |\mathbb{E}[\tilde{\mu}(w(S))]| - \epsilon \geq |\mathbb{E}[\tilde{\mu}(w(S^*))]| - 2\epsilon \geq |\mathbb{E}[\mu(w(S^*))]| - 3\epsilon$$

The proof is complete. \square

2.2 Approximating the Utility Function

In this subsection, we discuss the issue of approximating μ . In particular, we develop a generic algorithm that takes as a subroutine an algorithm AP for approximating functions in a bounded interval domain, and approximates $\mu(x)$ in the infinite domain $[0, +\infty)$. In the next subsection, we use the Fourier series expansion as the choice of AP and show that important classes of utility functions can be approximated well.

There are many works on approximating functions using short exponential sums, e.g., the Fourier decomposition approach [52], Prony's method [45], and many others [9, 10]. However, their approximations are done over a finite interval domain, say $[-\pi, \pi]$ or over a finite number of discrete points. No error bound can be guaranteed outside the domain. Our algorithm is a generic procedure that turns an algorithm that can approximate functions over $[-\pi, \pi]$ into one that can approximate our utility function μ over $[0, +\infty)$, by utilizing the fact that $\lim_{x \rightarrow \infty} \mu(x) = 0$.

Since $\lim_{x \rightarrow \infty} \mu(x) = 0$, for any ϵ , there exist a point T_ϵ such that $\mu(x) \leq \epsilon \quad \forall x > T_\epsilon$. Since we assume the utility function μ is specified as a part of the problem but not a part of the input instance, T_ϵ is a constant for any constant ϵ . We also assume there is an algorithm AP that, for any function f (under some conditions specified later), can produce an exponential sum $\hat{f}(x) = \sum_{i=1}^L c_i \phi_i^x$ which is an ϵ -approximation of $f(x)$ in $[-\pi, \pi]$ such that $|\phi_i| \leq 1$ and L depends only on ϵ and f . In fact, we can assume w.l.o.g. that AP can approximate $f(x)$ over $[-B, B]$ for any $B = O(1)$. This is because we can apply AP to the scaled version $g(x) = f(x \cdot \frac{B}{\pi})$ (which is defined on $[-\pi, \pi]$) and then scale the obtained approximation $\hat{g}(x)$ back to $[-B, B]$, i.e., the final approximation is $\hat{f}(x) = \hat{g}(\frac{\pi}{B} \cdot x)$. Scaling a function by a constant factor $\frac{B}{\pi}$ typically does not affect the smoothness of f in any essential way and we can still apply AP. Recall that our goal is to produce an exponential sum that is an ϵ -approximation for $\mu(x)$ in $[0, +\infty)$. We denote this procedure by ESUM.

Algorithm: ESUM

1. Initially, we slightly change function $\mu(x)$ to a new function $\hat{\mu}(x)$ as follows: We require $\hat{\mu}(x)$ is a “smooth” function in $[-2T_\epsilon, 2T_\epsilon]$ such that $\hat{\mu}(x) = \mu(x)$ for all $x \in [0, T_\epsilon]$; $\hat{\mu}(x) = 0$ for $|x| \geq 2T_\epsilon$. We choose $\hat{\mu}(x)$ in $[-2T_\epsilon, 0]$ and $[T_\epsilon, 2T_\epsilon]$ such that $\hat{\mu}(x)$ is smooth. We do not specify the exact smoothness requirements now since they may depend on the choice of AP. Note that there may be many ways to interpolate μ such that the above conditions are satisfied (see Example 1 below). The only properties we need are: (1) $\hat{\mu}$ is amenable to algorithm AP; (2) $|\hat{\mu}(x) - \mu(x)| \leq \epsilon \quad \forall x \geq 0$.
2. We apply AP to $f(x) = \eta^x \hat{\mu}(x)$ over domain $[-hT_\epsilon, hT_\epsilon]$ ($\eta \geq 1$ and $h \geq 2$ are constants to be determined later). Suppose the resulting exponential sum $\hat{f}(x) = \sum_{i=1}^L c_i \phi_i^x$ which is an ϵ -approximation of f on $[-hT_\epsilon, hT_\epsilon]$.
3. Let $\tilde{\mu}(x) = \sum_{i=1}^L c_i (\frac{\phi_i}{\eta})^x$, which is our final approximation of $\mu(x)$ on $[0, \infty)$.

Example 1 Consider the utility function $\mu(x) = 1/(x+1)$. Let $T_\epsilon = \frac{1}{\epsilon} - 1$. So $\mu(x) < \epsilon$ for all $x > T_\epsilon$. Now we create function $\hat{\mu}(x)$ according to the first step of ESUM. If we only require $\hat{\mu}(x)$ to be continuous, then we can use, for instance, the following piecewise function: $\hat{\mu}(x) = \frac{1}{x+1}, x \in [0, T_\epsilon]$; $\hat{\mu}(x) = \frac{x}{T_\epsilon} + \frac{2}{\epsilon}, x \in [T_\epsilon, 2T_\epsilon]$; $\hat{\mu}(x) = 0, x > 2T_\epsilon$; $\hat{\mu}(x) = -\hat{\mu}(x), x < 0$. It is easy to see that $\hat{\mu}$ is continuous and ϵ -approximates μ . \square

By setting $\eta = 2$ and

$$h \geq \frac{\log(\sum_{i=1}^L |c_i|/\epsilon)}{T_\epsilon}, \quad (3)$$

we can show the following theorem.

Lemma 3 $\tilde{\mu}(x)$ is a 2ϵ -approximation of $\mu(x)$.

Proof: We know that $|\hat{f}(x) - f(x)| \leq \epsilon$ for $x \in [0, hT_\epsilon]$. Therefore, we have that

$$|\tilde{\mu}(x) - \hat{\mu}(x)| = \left| \frac{\hat{f}(x)}{\eta^x} - \frac{f(x)}{\eta^x} \right| \leq \frac{\epsilon}{\eta^x} \leq \epsilon.$$

Combining with $|\hat{\mu}(x) - \mu(x)| \leq \epsilon$, we obtain $|\tilde{\mu}(x) - \mu(x)| \leq 2\epsilon$ for $x \in [0, hT_\epsilon]$. For $x > hT_\epsilon$, we can see that

$$|\tilde{\mu}(x)| = \left| \sum_{i=1}^L c_i \left(\frac{\phi_i}{\eta}\right)^x \right| \leq \sum_{i=1}^L |c_i| \left(\frac{\phi_i}{\eta}\right)^x \leq \frac{1}{2^x} \sum_{i=1}^L |c_i| \leq \frac{1}{2^{hT_\epsilon}} \sum_{i=1}^L |c_i| \leq \epsilon$$

Since $\mu(x) < \epsilon$ for $x > hT_\epsilon$, the proof is complete. \square

Remark: Since we do not know c_i before applying AP, we need to set h to be a constant (only depending on μ and ϵ) such that (3) is always satisfied. In particular, we need to provide an upper bound for $\sum_{i=1}^L |c_i|$. In the next subsection, we use the Fourier series decomposition as the choice for AP, which allows us to provide such a bound for a large class of functions.

2.3 A Particular Choice of AP: The Fourier Series Approach

Now, we discuss the choice of algorithm AP and the conditions that $f(x)$ needs to satisfy so that it is possible to approximate $f(x)$ by a short exponential sum in a bounded interval. In fact, if we know in advance that there is a short exponential sum that can approximate f , we can use the algorithms developed in [10, 11] (for continuous case) and [9] (for discrete case). However, those works do not provide an easy characterization of the class of functions. From now on, we restrict ourselves to the classic Fourier series technique, which has been studied extensively and allows such characterizations.

Consider the partial sum of the Fourier series of the function $f(x)$:

$$(S_N f)(x) = \sum_{k=-N}^N c_k e^{ikx}$$

where the Fourier coefficient $c_k = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(x) e^{-ikx} dx$. It has $L = 2N + 1$ terms. Since $f(x)$ is a real function, we have $c_k = c_{-k}$ and the partial sum is also real. We are interested in the question under which conditions does the function $S_N f$ converge to f (as N increases) and what is convergence rate? Roughly speaking, the more “smooth” f is, the faster $S_N f$ converges to f . In general, this question is extremely intricate and deep and is one of the central topics in the area of harmonic analysis. In the following, we give one classic result about the convergence of Fourier series and show how to use it in our problem. Then we provide a few concrete examples.

We say f satisfies the α -Hölder condition if $|f(x) - f(y)| \leq C |x - y|^\alpha$, for some constant C and $\alpha > 0$ and any x and y . The constant C is called the Hölder coefficient of f , also denoted as $|f|_{C^{0,\alpha}}$. We say f is C -Lipschitz if f satisfies 1-Hölder condition with coefficient C .

Example 2 It is easy to check that the utility function μ in Example 1 is 1-Lipschitz since $|\frac{d\mu(x)}{dx}| \leq 1$ for $x \geq 0$. We can also see that (1) is $\frac{1}{8}$ -Lipschitz.

We need the following classic result of Jackson.

Theorem 4 (See e.g., [47]) If f satisfies the α -Hölder condition, it holds that

$$|f(x) - (S_N f)(x)| \leq O\left(\frac{|f|_{C^{0,\alpha}} \ln N}{N^\alpha}\right).$$

For later development, we need a few simple lemmas. The proofs of these lemmas are straightforward and thus omitted here.

Lemma 4 Suppose $f : [a, c] \rightarrow \mathbb{R}$ is a continuous function which consists of two pieces $f_1 : [a, b] \rightarrow \mathbb{R}$ and $f_2 : [b, c] \rightarrow \mathbb{R}$. If both f_1 and f_2 satisfy the α -Hölder condition with Hölder coefficient C , then $|f|_{C^{0,\alpha}} \leq 2C$.

Lemma 5 Suppose $f : [a, c] \rightarrow \mathbb{R}$ is a continuous function satisfying the α -Hölder condition with Hölder coefficient C . Then, for $g(x) = f(hx)$ for some constant h , we have $|g|_{C^{0,\alpha}} \leq Ch^\alpha$.

Using Theorem 4 and Lemma 5, we obtain the following corollary.

Corollary 1 Suppose $f \in C^0[-hT_\epsilon, hT_\epsilon]$ satisfies the α -Hölder condition with $|f|_{C^{0,\alpha}} = O(1)$ and $N = O(hT_\epsilon(\frac{1}{\epsilon} \log \frac{1}{\epsilon})^{1/\alpha})$. Then, it holds that $|f(x) - (S_N f)(x)| \leq \epsilon$ for $x \in [-hT_\epsilon, hT_\epsilon]$.

Everything is in place to prove Theorem 2. Consider the algorithm AP. If μ is α -Hölder with coefficient $O(1)$, we can construct $\hat{\mu}$ which is also α -Hölder with coefficient $O(1)$, by Lemma 4. Then, we can easily see that $f(x) = \eta^x \hat{\mu}(x)$ is also α -Hölder with coefficient $O(1)$ in $[-hT_\epsilon, hT_\epsilon]$ for any $\eta = 2$. Hence, we can apply Corollary 1. By Lemma 3, we complete the proof of Theorem 2.

How to Choose h : Now, we discuss the issue left in Section 2.2, that is how to choose h (the value should be independent of c_i s and L) to satisfy (3), when μ satisfies the α -Hölder condition for some $\alpha > 1/2$. Indeed, we can choose $h = O(\frac{1}{T_\epsilon} \log \frac{1}{\epsilon})$. See Appendix D for the details.

3 Applications

We first consider two utility functions $\chi(x)$ and $\tilde{\chi}(x)$ presented in the introduction. Note that maximizing $\mathbb{E}[\chi(w(S))]$ is equivalent to maximizing $\Pr(w(S) \leq 1)$. The following lemma is straightforward.

Lemma 6 *For any solution S ,*

$$\Pr(w(S) \leq 1) \leq \mathbb{E}[\tilde{\chi}(w(S))] \leq \Pr(w(S) \leq 1 + \delta).$$

Corollary 2 *Suppose there is a pseudopolynomial time algorithm for the exact version of A . Then, for any fixed constants $\epsilon > 0$ and $\delta > 0$, there is an algorithm that runs in time $(\frac{n}{\epsilon})^{O(\frac{1}{\epsilon^2} \log \frac{1}{\epsilon})}$, and produces a solution $S \in \mathcal{F}$ such that*

$$\Pr(w(S) \leq 1 + \delta) + \epsilon \geq \max_{S' \in \mathcal{F}} \Pr(w(S') \leq 1)$$

Proof: By Theorem 1, Theorem 2 and Lemma 6, we can easily obtain the corollary. Note that we can choose $T_\epsilon = 2$ for any $\epsilon > 0$. Thus $h = O(\log \frac{1}{\epsilon})$ and $L = O(\frac{1}{\epsilon^2} \log \frac{1}{\epsilon})$. \square

Now, let us see some applications of our general results to specific problems.

Stochastic Shortest Path: Finding a path with the exact target length (we allow non-simple paths)⁵ can be easily done in pseudopolynomial time by dynamic programming. Therefore, as discussed in Section 1.1, Corollary 2 generalizes several results for stochastic shortest path in prior work [43, 41].

Stochastic Spanning Tree: Our objective is to find a spanning tree T in the given probabilistic graph such that $\Pr(w(T) \leq 1)$ is maximized. Polynomial time algorithms have been developed for Gaussian distributed edges [30, 23]. To the best of our knowledge, no approximation algorithm with provable guarantee is known for other distributions. Noticing there exists a pseudopolynomial time algorithm for the exact spanning tree problem [6], we can directly apply Corollary 2.

Stochastic k -Median on Trees: The problem asks for a set S of k nodes in the given probabilistic tree G such that $\Pr(\sum_{v \in V(G)} \text{dis}(v, S) \leq 1)$ is maximized, where $\text{dis}(v, S)$ is the minimum distance from v to any node in S in the tree metric. The k -median problem can be solved optimally in polynomial time on trees by dynamic programming [31]. In fact, we can easily modify the dynamic program to get a pseudopolynomial time algorithm for the exact version. We omit the details.

Stochastic Knapsack with Random Sizes: We are given a set U of n items. Each item i has a random size w_i and a deterministic profit v_i . We are also given a positive constant $0 \leq \gamma \leq 1$. The goal is to find a subset $S \subseteq U$ such that $\Pr(w(S) \leq 1) \geq \gamma$ and the total profit $v(S) = \sum_{i \in S} v_i$ is maximized.

⁵The exact version of *simple* path is NP-hard, since it includes the Hamiltonian path problem as a special case.

If the profits of the items are polynomially bounded integers, we can see the optimal profit is also a polynomially bounded integer. We can first guess the optimal profit. For each guess g , we solve the following problem: find a subset S of items such that the total profit of S is exactly g and $\mathbb{E}[\tilde{\chi}(w(S))]$ is maximized. The exact version of the deterministic problem is to find a solution S with a given total size and a given total profit, which can be easily solved in pseudopolynomial time by dynamic programming. Therefore, by Corollary 2, we can easily show that we can find in polynomial time a set S of items such that the total profit $v(S)$ is at least the optimum and $\Pr(w(S) \leq 1 + \epsilon) \geq (1 - \epsilon)\gamma$ for any constant ϵ and γ .

If the profits are general integers, we can use the standard scaling technique to get a $(1 - \epsilon)$ -approximation for the total profit. See Appendix E for the details. In sum, we have obtained the following result.

Theorem 5 *For any constants $\epsilon > 0$ and $\gamma > 0$, there is a polynomial time algorithm to compute a set S of items such that the total profit $v(S)$ is within a $1 - \epsilon$ factor of the optimum and $\Pr(w(S) \leq 1 + \epsilon) \geq (1 - \epsilon)\gamma$.*

Recently, Bhalgat et al. [13, Theorem 8.1] obtained the same result, with a running time $n^{2^{\text{poly}(1/\epsilon)}}$, while our running time is $(\frac{n}{\epsilon})^{O(\frac{1}{\epsilon^2} \log \frac{1}{\epsilon})} = n^{\text{poly}(1/\epsilon)}$.

Moreover, we can easily extend our algorithm to generalizations of the knapsack problem if the corresponding exact version has a pseudopolynomial time algorithm. For example, we can get the same result for the partial-ordered knapsack problem with tree constraints [22, 49]. In this problem, items must be chosen in accordance with specified precedence constraints and these precedence constraints form a partial order and the underlining undirected graph is a tree (or forest). A pseudopolynomial algorithm for this problem is presented in [49].

Stochastic Knapsack with Random Profits: We are given a set U of n items. Each item i has a deterministic size w_i and a random profit v_i . The goal is to find a subset of items that can be packed into a knapsack with capacity 1 and the probability that the profit is at least a given threshold T is maximized. Henig [29] and Carraway et al. [14] studied this problem for normally distributed profits and presented dynamic programming and branch and bound heuristics to solve this problem optimally.

We can solve the equivalent problem of minimizing the probability that the profit is at most the given threshold. It is straightforward to modify our algorithm to work for the minimization problem and we can also get an ϵ additive error for any $\epsilon > 0$. In fact, we can show that violation of the capacity constraint is necessary unless $P = NP$. See Appendix F for the details.

Theorem 6 *If the optimal probability is $\Omega(1)$, we can find in polynomial time a subset S of items such that $\Pr(v(S) > (1 - \epsilon)T) \geq (1 - \epsilon)OPT$ and $w(S) \leq 1 + \epsilon$, for any constant $\epsilon > 0$.*

4 Extensions

In this section, we discuss some extensions to our basic approximation scheme. We first consider optimizing a constant number of utility functions in Section 4.1. Then, we study the problem where the weight of each element is a random vector in Section 4.2.

4.1 Multiple Utility Functions

The problem we study in this section contains a set U of n elements. Each element e has a random weight w_e . We are also given d utility functions μ_1, \dots, μ_d and d positive numbers $\lambda_1, \dots, \lambda_d$. We assume d is a constant. A feasible solution consists of d subsets of elements that satisfy some property. Our objective is to find a feasible solution S_1, \dots, S_d such that $\mathbb{E}[\mu_i(w(S_i))] \geq \lambda_i$ for all $1 \leq i \leq d$.

We can easily extend our basic approximation scheme to the multiple utility functions case as follows. We decompose these utility functions into short exponential sums using ESUM as before. Then, for each utility function, we maintain $(n/\epsilon)^{O(L)}$ configurations. Therefore, we have $(n/\epsilon)^{O(dL)}$ configurations in total and we would like to compute the values for these configurations. We denote the deterministic version of the problem under consideration by A . The exact version of A asks for a feasible solution S_1, \dots, S_d such that the total weight of S_i is exactly the given number t_i for all i . Following an argument similar to Lemma 2, we can easily get the following generalization of Theorem 1.

Theorem 7 *Assume that there is a pseudopolynomial algorithm for the exact version of A . Further assume that given any $\epsilon > 0$, we can ϵ -approximate each utility function by an exponential sum with at most L terms. Then, there is an algorithm that runs in time $(n/\epsilon)^{O(dL)}$ and finds a feasible solution S_1, \dots, S_d such that $\mathbb{E}[\mu_i(w(S_i))] \geq \lambda_i - \epsilon$ for $1 \leq i \leq d$, if there is a feasible solution for the original problem.*

Now let us consider two simple applications of the above theorem.

Stochastic Multiple Knapsack: In this problem we are given a set U of n items, d knapsacks with capacity 1, and d constants $0 \leq \gamma_i \leq 1$. We assume d is a constant. Each item i has a random size w_i and a deterministic profit v_i . Our objective is to find d disjoint subsets S_1, \dots, S_d such that $\Pr(w(S_i) \leq 1) \geq \gamma_i$ for all $1 \leq i \leq d$ and $\sum_{i=1}^d v(S_i)$ is maximized. The exact version of the problem is to find a packing such that the load of each knapsack i is *exactly* the given value t_i . It is not hard to show this problem can be solved in pseudopolynomial time by standard dynamic programming. If the profits are general integers, we also need the scaling technique as in stochastic knapsack with random sizes (Appendix E). In sum, we can get the following generalization of Theorem 5.

Theorem 8 *For any constants $d \in \mathbb{N}$, $\epsilon > 0$ and $0 \leq \gamma_i \leq 1$ for $1 \leq i \leq d$, there is a polynomial time algorithm to compute d disjoint subsets S_1, \dots, S_d such that the total profit $\sum_{i=1}^d v(S_i)$ is within a $1 - \epsilon$ factor of the optimum and $\Pr(w(S_i) \leq 1 + \epsilon) \geq (1 - \epsilon)\gamma_i$ for $1 \leq i \leq d$.*

Stochastic Multidimensional Knapsack: In this problem we are given a set U of n items and a constant $0 \leq \gamma \leq 1$. Each item i has a deterministic profit v_i and a random size which is a random d -dimensional vector $\mathbf{w}_i = \{w_{i1}, \dots, w_{id}\}$. We assume d is a constant. Our objective is to find a subset S of items such that $\Pr(\bigwedge_{j=1}^d (\sum_{i \in S} w_{ij} \leq 1)) \geq \gamma$ and total profit is maximized. This problem can be also thought as the fixed set version of the stochastic packing problem considered in [19, 13]. We first assume the components of each size vector are independent. The correlated case will be addressed in the next subsection.

For ease of presentation, we assume $d = 2$ from now on. Extension to general constant d is straightforward. We can solve the problem by casting it into a multiple utility problem as follows. For each item i , we create two copies i_1 and i_2 . The copy i_j has a random weight w_{ij} . A feasible solution consists of two sets S_1 and S_2 such that S_1 (S_2) only contains the first (second) copies of the elements and S_1 and S_2 correspond to exactly the same subset of original elements. We enumerate all such pairs (γ_1, γ_2) such that $\gamma_1 \gamma_2 \geq \gamma$ and $\gamma_i \in [\gamma, 1]$ is a power of $1 - \epsilon$ for $i = 1, 2$. Clearly, there are a polynomial number of such pairs. For each pair (γ_1, γ_2) , we solve the following problem: find a feasible solution S_1, S_2 such that $\Pr(\sum_{i \in S_j} w_{ij} \leq 1) \geq \gamma_j$ for all $j = 1, 2$ and total profit is maximized. Using the scaling technique and Theorem 7 for optimizing multiple utility functions, we can get a $(1 - \epsilon)$ -approximation for the optimal profit and $\Pr(\bigwedge_{j=1}^2 (\sum_{i \in S_j} w_{ij} \leq 1)) = \prod_{j=1}^2 \Pr(\sum_{i \in S_j} w_{ij} \leq 1) \geq (1 - O(\epsilon))\gamma_1 \gamma_2 \geq (1 - O(\epsilon))\gamma$.

We note that the same result for independent components can be also obtained by using the discretization technique developed for the adaptive version of the problem in [13]⁶. If the components of each size vector

⁶With some changes of the discretization technique, the correlated case can be also handled [12].

are correlated, we can not decompose the problem into two 1-dimensional utilities as in the independent case. Now, we introduce a new technique to handle the correlated case.

4.2 Multidimensional Weight

The general problem we study contains a set U of n elements. Each element e has a random weight vector $w_i = (w_{i1}, \dots, w_{id})$. We assume d is a constant. We are also given a utility functions $\mu : \mathbb{R}^d \rightarrow \mathbb{R}^+$. A feasible solution is a subset of elements satisfying some property. We use $w(S)$ as a shorthand notation for vector $(\sum_{i \in S} w_{i1}, \dots, \sum_{i \in S} w_{id})$. Our objective is to find a feasible solution S such that $\mathbb{E}[\mu(w(S))]$ is maximized.

From now on, x and k denote d -dimensional vectors and kx (or $k \cdot x$) denotes the inner product of k and x . As before, we assume $\mu(x) \in [0, 1]$ for all $x \geq 0$ and $\lim_{|x| \rightarrow +\infty} \mu(x) = 0$, where $|x| = \max(x_1, \dots, x_d)$. Our algorithm is almost the same as in the one dimension case and we briefly sketch it here. We first notice that expected utilities decompose for exponential utility functions, i.e., $\mathbb{E}[\phi^{k \cdot w(S)}] = \prod_{i \in S} \mathbb{E}[\phi^{k \cdot w_i}]$. Then, we attempt to ϵ -approximate the utility function $\mu(x)$ by a short exponential sum $\sum_{|k| \leq N} c_k \phi_k^{kx}$ (there are $O(N^d)$ terms). If this can be done, $\mathbb{E}[\phi^{k \cdot w(S)}]$ can be approximated by $\sum_{|k| \leq N} c_k \mathbb{E}[\phi^{k \cdot w(S)}]$. Using the same argument as in Theorem 1, we can show that there is a polynomial time algorithm that can find a feasible solution S with $\mathbb{E}[\mu(w(S))] \geq OPT - \epsilon$ for any $\epsilon > 0$, provided that a pseudopolynomial algorithm exists for the exact version of the deterministic problem.

To approximate the utility function $\mu(x)$, we need the multidimensional Fourier series expansion of a function $f : \mathbb{C}^d \rightarrow \mathbb{C}$ (assuming f is 2π -periodic in each axis): $f(x) \sim \sum_{k \in \mathbb{Z}^d} c_k e^{ikx}$ where $c_k = \frac{1}{(2\pi)^d} \int_{x \in [-\pi, \pi]^d} f(x) e^{-ikx} dx$. The *rectangular partial sum* is defined to be

$$S_N f(x) = \sum_{|k_1| \leq N} \dots \sum_{|k_d| \leq N} c_k e^{ikx}.$$

It is known that the rectangular partial sum $S_N f(x)$ converges uniformly to $f(x)$ in $[-\pi, \pi]^d$ for many function classes as n tends to infinity. In fact, a generalization of Theorem 4 to $[-\pi, \pi]^d$ also holds [4]: If f satisfies the α -Hölder condition, then

$$|f(x) - (S_N f)(x)| \leq O\left(\frac{|f|_{C^{0,\alpha}} \ln^d N}{N^\alpha}\right) \quad \text{for } x \in [-\pi, \pi]^d.$$

Now, we have an algorithm AP that can approximate a function in a bounded domain. It is also straightforward to extend ESUM to the multidimensional case. Hence, we can ϵ -approximate μ by a short exponential sum in $[0, +\infty)^d$, thereby proving the multidimensional generalization of Theorem 2. Let us consider an application of our result.

Stochastic Multidimensional Knapsack (Revisited): We consider the case where the components of each weight vector can be correlated. Note that the utility function χ_2 corresponding to this problem is the two dimensional threshold function: $\chi_2(x, y) = 1$ if $x \leq 1$ and $y \leq 1$; $\chi_2(x, y) = 0$ otherwise. As in the one dimensional case, we need to consider a continuous version $\tilde{\chi}_2$ of χ_2 (see Figure 1(3)). By the result in this section and a generalization of Lemma 6 to higher dimension, we can get the following.

Theorem 9 *For any constants $d \in \mathbb{N}$, $\epsilon > 0$ and $0 \leq \gamma \leq 1$, there is a polynomial time algorithm for finding a set S of items such that the total profit $v(S)$ is $1 - \epsilon$ factor of the optimum and $\Pr(\bigwedge_{j=1}^d (\sum_{i \in S} w_{ij} \leq 1 + \epsilon)) \geq (1 - \epsilon)\gamma$.*

5 Discussions

Convergence of Fourier series: The convergence of the Fourier series of a function is a classic topic in harmonic analysis. Whether the Fourier series converges to the given function and the rate of the convergence typically depends on a variety of smoothness condition of the function. We refer the readers to [52] for a more comprehensive treatment of this topic. We note that we could obtain a smoother version of χ (e.g., see Figure 1(2)), instead of the piecewise linear $\tilde{\chi}$, and then use Theorem 4 to obtain a better bound for L . This would result in an even better running time. Our choice is simply for the ease of presentation.

Discontinuous utility functions: If the utility function μ is discontinuous, e.g., the threshold function, then the partial Fourier series behaves poorly around the discontinuity (this is known as the *Gibbs phenomenon*). However, informally speaking, as the number of Fourier terms increases, the poorly-behaved strip around the edge becomes narrower. Therefore, if the majority of the probability mass of our solution lies outside the strip, we can still guarantee a good approximation of the expected utility. There are also techniques to reduce the effects of the Gibbs phenomenon (See e.g., [25]). We leave the problem of directly dealing with discontinuous utility functions, especially the threshold function, to obtain a true approximation (instead of a bi-criterion approximation) as an interesting open problem.

6 Conclusion

We consider the problem of maximizing expected utility for many stochastic combinatorial problems, such as shortest path, spanning tree and knapsack. We develop a polynomial time approximation scheme with additive error ϵ for any $\epsilon > 0$. A key ingredient in our algorithm is to decompose the utility function into a short exponential sum. In this paper, we use the Fourier series technique to fulfill the task. Exploring other decomposition approaches is an interesting future work. Our general approximation framework may be useful for other stochastic optimization problems. One major open problem is to obtain approximations with reasonable multiplicative factors, or nontrivial inapproximability results, for the utility maximization problem.

7 Acknowledgments

We would like to thank Evdokia Nikolova for providing an extended version of [43] and many helpful discussions. We also would like to thank Chandra Chekuri for pointing to us the recent work [13] and Anand Bhalgat for some clarifications of the same work. This work was supported in part by NSF Grant IIS-0916736, a DARPA grant, National Basic Research Program of China Grants 2007CB807900, 2007CB807901, and the National Natural Science Foundation of China Grants 61033001, 61061130540, 61073174.

References

- [1] St. petersburg paradox. http://en.wikipedia.org/wiki/St._Petersburg_paradox.
- [2] H. Ackermann, A. Newman, H. Röglin, and B. Vöcking. Decision making based on approximate and smoothed pareto curves. *Algorithms and Computation*, pages 675–684, 2005.
- [3] S. Agrawal, A. Saberi, and Y. Ye. Stochastic Combinatorial Optimization under Probabilistic Constraints. *Arxiv preprint arXiv:0809.0460*, 2008.

- [4] S. Alimov, R. Ashurov, and A. Pulatov. Multiple fourier series and fourier integrals, in commutative harmonic analysis. IV: Harmonic analysis in \mathbb{R}_n . *Encyclopedia of Mathematical Science*, 42, 1992.
- [5] N. Bansal, A. Gupta, J. Li, J. Mestre, V. Nagarajan, and A. Rudra. When LP is the Cure for Your Matching Woes: Improved Bounds for Stochastic Matchings. *European Symposium on Algorithms*, pages 218–229, 2010.
- [6] F. Barahona and W. Pulleyblank. Exact arborescences, matchings and cycles. *Discrete Applied Mathematics*, 16(2):91–99, 1987.
- [7] J. Bard and J. Bennett. Arc reduction and path preference in stochastic acyclic networks. *Management Science*, 37(2):198–215, 1991.
- [8] D. Bernoulli. Exposition of a new theory on the measurement of risk. *Econometrica*, 22(1):22–36, 1954. Originally published in 1738; translated by Dr. Lousie Sommer.
- [9] G. Beylkin and L. Monzón. On Generalized Gaussian Quadratures for Exponentials and Their Applications* 1. *Applied and Computational Harmonic Analysis*, 12(3):332–373, 2002.
- [10] G. Beylkin and L. Monzón. On approximation of functions by exponential sums. *Applied and Computational Harmonic Analysis*, 19(1):17–48, 2005.
- [11] G. Beylkin and L. Monzón. Approximation by exponential sums revisited. *Applied and Computational Harmonic Analysis*, 28(2):131–149, 2010.
- [12] A. Bhargat, 2011. Personal Communication.
- [13] A. Bhargat, A. Goel, and S. Khanna. Improved approximation results for stochastic knapsack problems. In *ACM-SIAM symposium on Discrete algorithms*, 2011.
- [14] R. Carraway, R. Schmidt, and L. Weatherford. An algorithm for maximizing target achievement in the stochastic knapsack problem with normal returns. *Naval research logistics*, 40(2):161–173, 1993.
- [15] C. Chekuri and S. Khanna. A PTAS for the multiple knapsack problem. In *ACM-SIAM symposium on Discrete algorithms*, pages 213–222, 2000.
- [16] N. Chen, N. Immorlica, A. Karlin, M. Mahdian, and A. Rudra. Approximating matches made in heaven. *International Colloquium on Automata, Languages and Programming*, pages 266–278, 2009.
- [17] W. Cheney and W. Light. *A Course in Approximation Theory*. Brook/Cole Publishing Company, 2000.
- [18] R. Cheng, J. Chen, and X. Xie. Cleaning uncertain data with quality guarantees. *Proceedings of the VLDB Endowment*, 1(1):722–735, 2008.
- [19] B. Dean, M. Goemans, and J. Vondrák. Adaptivity and approximation for stochastic packing problems. In *ACM-SIAM symposium on Discrete algorithms*, pages 395–404, 2005.
- [20] B. Dean, M. Goemans, and J. Vondrak. Approximating the Stochastic Knapsack Problem: The Benefit of Adaptivity. *Mathematics of Operations Research*, 33(4):945, 2008.
- [21] P. Fishburn. *Utility Theory and Decision Making*. John Wiley & Sons, Inc, 1970.

- [22] M. Garey and D. Johnson. “*Computers and Intractability: A Guide to the Theory of NP-Completeness*”. W.H. Freeman, 1979.
- [23] S. Geetha and K. Nair. On stochastic spanning tree problem. *Networks*, 23(8):675–679, 1993.
- [24] A. Goel and P. Indyk. Stochastic load balancing and related problems. In *Annual Symposium on Foundations of Computer Science*, page 579, 1999.
- [25] D. Gottlieb and C. Shu. On the Gibbs phenomenon and its resolution. *SIAM review*, 39(4):644–668, 1997.
- [26] V. Goyal and R. Ravi. Chance constrained knapsack problem with random item sizes. *To appear in Operation Research Letter*, 2009.
- [27] S. Guha and K. Munagala. Adaptive Uncertainty Resolution in Bayesian Combinatorial Optimization Problems. *To appear in ACM Transactions on Algorithms*, 2008.
- [28] A. Gupta, M. Pál, R. Ravi, and A. Sinha. Boosted sampling: approximation algorithms for stochastic optimization. In *ACM Symposium on Theory of Computing*, pages 417–426. ACM, 2004.
- [29] M. Henig. Risk criteria in a stochastic knapsack problem. *Operations Research*, 38(5):820–825, 1990.
- [30] H. Ishii, S. Shiode, and T. Nishida Yoshikazu. Stochastic spanning tree problem. *Discrete Applied Mathematics*, 3(4):263–273, 1981.
- [31] O. Kariv and S. Hakimi. An algorithmic approach to network location problems. II: The p-medians. *SIAM Journal on Applied Mathematics*, 37(3):539–560, 1979.
- [32] J. Kleinberg, Y. Rabani, and É. Tardos. Allocating bandwidth for bursty connections. In *ACM Symposium on Theory of Computing*, page 673, 1997.
- [33] J. Li and A. Deshpande. Consensus answers for queries over probabilistic databases. In *ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, 2009.
- [34] J. Li and A. Deshpande. Ranking continuous probabilistic datasets. *Proceedings of the VLDB Endowment*, 3(1), 2010.
- [35] J. Li, B. Saha, and A. Deshpande. A unified approach to ranking in probabilistic databases. In *Proceedings of the VLDB Endowment*, 2009.
- [36] R. Loui. Optimal paths in graphs with stochastic or multidimensional weights. *Communications of the ACM*, 26(9):670–676, 1983.
- [37] R. Martin. The St. Petersburg Paradox. *The Stanford Encyclopedia of Philosophy*, 2004. <http://plato.stanford.edu/archives/fall12004/entries/paradox-stpetersburg>.
- [38] S. Mittal and A. Schulz. A general framework for designing approximation schemes for combinatorial optimization problems with many objectives combined into one. *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques*, pages 179–192, 2008.
- [39] I. Murthy and S. Sarkar. Exact algorithms for the stochastic shortest path problem with a decreasing deadline utility function. *European Journal of Operational Research*, 103(1):209–229, 1997.

- [40] I. Murthy and S. Sarkar. Stochastic shortest path problems with piecewise-linear concave utility functions. *Management Science*, 44(11):125–136, 1998.
- [41] E. Nikolova. Approximation Algorithms for Reliable Stochastic Combinatorial Optimization. *International Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, pages 338–351, 2010.
- [42] E. Nikolova, M. Brand, and D. Karger. Optimal route planning under uncertainty. In *Proceedings of International Conference on Automated Planning and Scheduling*, 2006.
- [43] E. Nikolova, J. Kelner, M. Brand, and M. Mitzenmacher. Stochastic shortest paths via quasi-convex maximization. In *European Symposium on Algorithms*, pages 552–563, 2006.
- [44] F. Oberhettinger. *Fourier transforms of distributions and their inverses: a collection of tables*. Academic press, 1973.
- [45] M. R. Osborne and G. K. Smyth. A modified prony algorithm for fitting sums of exponential functions. *SIAM Journal of Scientific Computing*, 1995.
- [46] C. Papadimitriou and M. Yannakakis. On the approximability of trade-offs and optimal access of web sources. In *Annual Symposium on Foundations of Computer Science*, 2000.
- [47] M. J. D. Powell. *Approximation theory and methods*. Cambridge University Press, 1981.
- [48] A. Ralston and R. Rabinowitz. *A First Course in Numerical Analysis*. 2001.
- [49] H. Safer, J. B. Orlin, and M. Dror. Fully polynomial approximation in multi-criteria combinatorial optimization, 2004. MIT Working Paper.
- [50] D. Shmoys and C. Swamy. An approximation scheme for stochastic linear programming and its application to stochastic integer programs. *J. ACM*, 53(6):1012, 2006.
- [51] C. Sigal, A. Pritsker, and J. Solberg. The stochastic shortest route problem. *Operations Research*, 28(5):1122–1129, 1980.
- [52] E. Stein and R. Shakarchi. *Fourier analysis: an introduction*. Princeton University Press, 2003.
- [53] C. Swamy. Risk-Averse Stochastic Optimization: Probabilistically-Constrained Models and Algorithms for Black-Box Distributions. *ACM-SIAM symposium on Discrete algorithms*, 2010.
- [54] C. Swamy and D. Shmoys. Approximation algorithms for 2-stage stochastic optimization problems. *ACM SIGACT News*, 37(1):46, 2006.
- [55] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton Univ. Press, 2nd edition, 1947.

A St. Petersburg Paradox

In this section, we briefly describe the St. Petersburg paradox. The paradox is named from Daniel Bernoulli's presentation of the problem, published in 1738 in the Commentaries of the Imperial Academy of Science of Saint Petersburg. Consider the following game: you pay a fixed fee X to enter the game. In the game, a fair coin is tossed repeatedly until a tail appears ending the game. The payoff of the game is 2^k where k is the number of heads that appears., i.e., you win 1 dollar if a tail appears on the first toss, 2 dollars if a head appears on the first toss and a tail on the second, 4 dollars if a head appears on the first two tosses and a tail on the third and so on. The question is what would be a fair fee X to enter the game? First, it is easy to see that the expected payoff is

$$\mathbb{E}[\text{payoff}] = \frac{1}{2} \cdot 1 + \frac{1}{4} \cdot 2 + \frac{1}{8} \cdot 4 + \frac{1}{16} \cdot 8 + \cdots = \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + \cdots = \sum_{k=1}^{\infty} \frac{1}{2} = \infty$$

If we use the expected payoff as a criterion for decision making, we should therefore play the game at any finite price X (no matter how large X is) since the expected payoff is always larger. However, researchers have done extensive survey and found that not many people would pay even 25 dollars to play the game [37], which significantly deviates from what the expected value criterion predicts. In fact, the paradox can be resolved by the expected utility theory with a logarithmic utility function, suggested by Bernoulli himself. We refer interested reader to [37, 1] for more information.

B Missing Proofs

Proof of Lemma 1: We first notice that

$$\mathbb{E}[\tilde{\mu}(w(S))] = \mathbb{E}\left[\sum_{k=1}^L c_k \phi_k^{w(S)}\right] = \sum_{k=1}^L c_k \mathbb{E}[\phi_k^{w(S)}].$$

Therefore, it suffices to show that for all $k = 1, \dots, L$,

$$|\mathbb{E}[\phi_k^{w(S)}] - e^{-\alpha_k \gamma + i\beta_k \delta}| \leq O\left(\frac{\epsilon}{L}\right).$$

First, we can see that

$$\arg(\mathbb{E}[\phi_k^{w(S)}]) - \beta_k \delta = \sum_{e \in S} (\arg(\mathbb{E}[\phi_k^{w_e}]) - b_k(e) \delta) \leq \sum_{e \in S} \delta \leq n \delta = \frac{\epsilon}{L}.$$

If $\sum_{e \in S} a_k(e) > J$, we know that

$$-\ln(|\mathbb{E}[\phi_k^{w(S)}]|) = \sum_{e \in S} (-\ln(|\mathbb{E}[\phi_k^{w_e}]|)) > J\gamma.$$

In this case, we have $\alpha_k = J$. Thus, we have

$$\left| |\mathbb{E}[\phi_k^{w(S)}]| - |e^{-\alpha_k \gamma}| \right| < e^{-J\gamma} = e^{\gamma \lceil \frac{\ln(\epsilon/L)}{\gamma} \rceil} < \frac{\epsilon}{L}.$$

If $\sum_{e \in S} a_k(e) \leq J$, we can see that

$$-\ln(|\mathbb{E}[\phi^{w(S)}]|) - \alpha_k \gamma = \sum_{e \in S} (-\ln(|\mathbb{E}[\phi^{w_e}]|) - \alpha_k(e) \gamma) \leq \sum_{e \in S} \gamma \leq n \gamma \leq \frac{\epsilon}{L}.$$

Since the derivative of e^x is less than 1 for $x < 0$, we can get

$$\left| |\mathbb{E}[\phi_k^{w(S)}]| - |e^{-\alpha_k \gamma}| \right| \leq |e^{-\alpha_k \gamma - \frac{\epsilon}{L}} - e^{-\alpha_k \gamma}| \leq \frac{\epsilon}{L}.$$

For any two complex numbers a, b with $|a| \leq 1$ and $|b| \leq 1$, if $||a| - |b|| < h$ and $|\arg(a) - \arg(b)| < h$, we can easily show that $|a - b| < O(h)$. The proof is complete. \square

Proof of Lemma 2: For each element e , we associate a new vector $\bar{\mathbf{a}}_e = \langle \bar{a}_1, \bar{b}_1, \dots, \bar{a}_L, \bar{b}_L \rangle$. If $a_i(e) > J$, we let $\bar{a}_i(e) = n(J+1)$ and $\bar{a}_i(e) = a_i(e)$ otherwise. Let $\bar{b}_i(e) = b_i(e)$ for all e and i . For each node v and each vector $\mathbf{a} = \langle \alpha_1, \beta_1, \dots, \alpha_L, \beta_L \rangle$ such that $0 \leq \alpha_i \leq n^2(J+1) \forall i$ and $0 \leq \beta_i \leq K \forall i$, we want to compute the value $\bar{\sigma}_v(\mathbf{a})$ which is defined as follows: $\bar{\sigma}_v(\mathbf{a}) = 1$ if and only if there is a feasible solution $S \in \mathcal{F}$ such that for all $j = 1, \dots, L$, $\beta_j = \sum_{e \in S} \bar{b}_j(e)$, and $\alpha_j = \sum_{e \in S} \bar{a}_j(e)$ (or more compactly, $\mathbf{a} = \sum_{e \in S} \bar{\mathbf{a}}_e$); $\bar{\sigma}_v(\mathbf{a}) = 0$ otherwise.

We can encode each vector as a nonnegative integer upper bounded by $(n^2 JK)^L = (\frac{n}{\epsilon})^{O(L)}$. Then, determining the value of a configuration is equivalent to determining whether there is a feasible solution S such that the total weight of S is exactly a given value. Suppose the pseudopolynomial time algorithm for the exact version of A runs in time $P_A(n, t)$ for some polynomial P_A . Therefore, the value of each such $\bar{\sigma}_v(\mathbf{a})$ can be also computed in time $P_A(n, (\frac{n}{\epsilon})^{O(L)}) = (\frac{n}{\epsilon})^{O(L)}$. Since J and K are bounded by $(\frac{n}{\epsilon})^{O(1)}$, the number of configuration is $(\frac{n}{\epsilon})^{O(L)}$. The value of $\sigma(\langle \alpha_1, \beta_1, \dots, \alpha_L, \beta_L \rangle)$ can be easily answered from the values of $\bar{\sigma}_v$ as follows :

1. If $\alpha_i < J \forall i$, $\sigma_v(\mathbf{a}) = \bar{\sigma}_v(\mathbf{a})$;
2. Denote $\mathbf{a}' = \langle \alpha'_1, \beta'_1, \dots, \alpha'_L, \beta'_L \rangle$ and $S = \{i \mid \alpha_i = J\}$. $\sigma_v(\mathbf{a}) = \max_{\mathbf{a}'} (\bar{\sigma}_v(\mathbf{a}') \mid \beta'_i = \beta_i \forall i, \alpha'_i \geq J \forall i \in S, \alpha'_i = \alpha_i \forall i \notin S)$.

The total running time is $(\frac{n}{\epsilon})^{O(L)} \times (\frac{n}{\epsilon})^{O(L)} = (\frac{n}{\epsilon})^{O(L)}$. \square

C Computing $\mathbb{E}[\alpha^{w_e}]$

If X is a random variable, then the *characteristic function* of X is defined as

$$G(z) = \mathbb{E}[e^{izX}].$$

We can see $\mathbb{E}[\alpha^{w_e}]$ is nothing but the value of the characteristic function of w_e evaluated at $-i \ln \alpha$ (here \ln is the complex logarithm function). For many important distributions, including negative binomial, Poisson, exponential, Gaussian, Chi-square and Gamma, a closed-form characteristic function is known. See [44] for a more comprehensive list.

Example 3 Consider the Poisson distributed w_e with mean λ , i.e., $\Pr(w_e = k) = \lambda^k e^{-\lambda} / k!$. Its characteristic function is known to be $G(z) = e^{\lambda(e^{iz} - 1)}$. Therefore,

$$\mathbb{E}[\alpha^{w_e}] = G(-i \ln \alpha) = e^{\lambda(\alpha - 1)}.$$

Example 4 For Gaussian distribution $N(\mu, \sigma^2)$, we know its characteristic function is $G(z) = e^{iz\mu - \frac{1}{2}\sigma^2 z^2}$. Therefore,

$$\mathbb{E}[\alpha^{w_e}] = G(-i \ln \alpha) = \alpha^{u + \frac{1}{2}\sigma^2 \ln \alpha}.$$

For some continuous distributions, no closed-form characteristic function is known and we need proper numerical approximation method.

If the support of the distribution is bounded, we can use for example Gauss-Legendre quadrature [48]. If the support is infinite, we can truncate the distribution and approximate the integral over the remaining finite interval; Generally speaking a quadrature method approximates $\int_a^b f(x)dx$ by a linear sum $\sum_{i=1}^k c_i f(x_i)$ where c_i and x_i are some constants independent of the function f . A typical practice is to use *composite rule*, that is to partition $[a, b]$ into N subintervals and approximate the integral using some quadrature formula over each subinterval. For the example of Gauss-Laguerre quadrature, assuming continuity of the $2k$ th derivative of $f(x)$ for some constant k , if we partition $[a, b]$ into M subintervals and apply Gauss-Legendre quadrature of degree k to each subinterval, the approximation error is

$$Error = \frac{(b-a)^{2k+1}}{M^{2k}} \frac{(k!)^4}{(2k+1)[(2k)!]^3} f^{(2k)}(\xi)$$

where ξ is some point in (a, b) [48, pp.116]. Let $\Delta = \frac{b-a}{M}$. If we treat k as a constant, the behavior of the error (in terms of Δ) is $Error(\Delta) = O(\Delta^{2k} \max_{\xi} f^{(2k)}(\xi))$. Therefore, if the support and $\max_{\xi} f^{(2k)}(\xi)$ are bounded by a polynomial, we can approximate the integral, in polynomial time, such that the error is $O(1/n^\beta)$ for any fixed integer β .

The next lemma shows that we do not lose too much even though we can only get an approximation of $\mathbb{E}[\alpha^{w_e}]$.

Lemma 7 Suppose in Theorem 3, we can only compute an approximate value of $\mathbb{E}[\phi_i^{w_e}]$, denoted by $E_{e,i}$, for each e and i , such that $|\mathbb{E}[\phi_i^{w_e}] - E_{e,i}| \leq O(n^{-\beta})$ for some positive integer β . Denote $E(S) = \sum_{k=1}^L c_k \prod_{e \in S} E_{e,i}$. For any solution S , we have that

$$|\mathbb{E}[\tilde{\mu}(w(S))] - E(S)| \leq O(n^{1-\beta}).$$

Proof: We need the following simple result (see [34] for a proof): a_1, \dots, a_n and e_1, \dots, e_n are complex numbers such that $|a_i| \leq 1$ and $|e_i| \leq n^{-\beta}$ for all i and some $\beta > 1$. Then, we have

$$\left| \prod_{i=1}^n (a_i + e_i) - \prod_{i=1}^n a_i \right| \leq O(n^{1-\beta}).$$

Since $|\phi_i| \leq 1$, we can see that

$$|\mathbb{E}[\phi_i^{w_e}]| = \left| \int_{x \geq 0} \phi_i^x p_e(x) dx \right| \leq 1.$$

The lemma simply follows by applying the above result and noticing that L and all c_k s are constants. \square

We can show that Theorem 1 still holds even though we only have the approximations of the $\mathbb{E}[\alpha^{w_e}]$ values. The proof is straightforward and omitted.

D How to Choose h

In this section, we discuss the issue left in Section 2.2, that is how to choose h (the value should be independent of c_i s and L) to satisfy (3), when μ satisfies the α -Hölder condition for some $\alpha > 1/2$.

We need the following results about the absolute convergence of Fourier coefficients. If f satisfies the α -Hölder condition for some $\alpha > 1/2$, then $\sum_{i=-\infty}^{+\infty} |c_i| \leq |f|_{C^{0,\alpha}} \cdot c_\alpha$ where c_α only depends on α [52].

Suppose the original utility function μ satisfies the α -Hölder condition with coefficient C , for some $\alpha > 1/2$. Now, we apply ESUM to μ . By Lemma 4, we know that the piecewise function μ satisfies α -Hölder condition with coefficient $2C$. Therefore, we can easily see that $f(x) = \hat{\mu}(x)\eta^x$ satisfies α -Hölder condition with coefficient at most $2^{1+2T_\epsilon}C$ on $[-hT_\epsilon, hT_\epsilon]$ (This is because $\hat{\mu}$ is non-zero only in $[-2T_\epsilon, 2T_\epsilon]$). According to Lemma 5, we have $|f(x \cdot \frac{hT_\epsilon}{2\pi})|_{C^{0,\alpha}} \leq 2^{1+2T_\epsilon}(\frac{hT_\epsilon}{2\pi})^\alpha C$.

Therefore, it is sufficient to set value h such that

$$hT_\epsilon \geq \log \frac{2^{1+2T_\epsilon}(\frac{hT_\epsilon}{2\pi})^\alpha C c_\alpha}{\epsilon} = 2T_\epsilon + O(\log(\frac{hT_\epsilon}{\epsilon})).$$

We can easily verify that we can satisfy the above condition by letting $h = O(\frac{1}{T_\epsilon} \log \frac{1}{\epsilon})$.

E Details for Stochastic Knapsack with Random Sizes

We first make a guess of the optimal profit, rounded down to the nearest power of $(1 + \epsilon)$. There are at most $\log_{1+\epsilon} \frac{n \max_i v_i}{\min_i v_i}$ guesses. For each guess g , we solve the following problem. We discard all items with a profit larger than g . Let $\Delta = \frac{\epsilon g}{n^2}$. For each item with a profit smaller than $\frac{\epsilon g}{n}$, we set its new profit to be $\bar{v}_i = 0$. Then, we scale each of the rest profits v_i to $\bar{v}_i = \Delta \lfloor \frac{v_i}{\Delta} \rfloor$. Now, we define the feasible set

$$\mathcal{F}(g) = \{S \mid \sum_{i \in S} (1 - 2\epsilon)g \leq \sum_{i \in S} \bar{v}_i \leq (1 + 2\epsilon)g\}.$$

Since there are at most $\frac{n^2}{\epsilon}$ distinct \bar{v} values, we can easily show that finding a solution S in $\mathcal{F}(g)$ with a given total size can be solved in pseudopolynomial time by dynamic programming.

Denote the optimal solution by S^* and the optimal profit by OPT . Suppose g is the right guess, i.e., $(\frac{1}{1+\epsilon})OPT \leq g \leq OPT$. We can easily see that for any solution S , we have that

$$(1 - \frac{1}{n}) \sum_{i \in S} v_i - \epsilon g \leq \sum_{i \in S} \bar{v}_i \leq \sum_{i \in S} v_i$$

where the first inequalities are due to $v_i \geq \frac{\epsilon g}{n}$ and we set at most ϵg profit to zero. Therefore, we can see $S^* \in \mathcal{F}(g)$. Applying Corollary 2, we obtain a solution S such that $\Pr(w(S) \leq 1 + \delta) + \epsilon \geq \Pr(w(S^*) \leq 1 + \delta)$. Moreover, the profit of this solution $v(S) = \sum_{i \in S} v_i \geq \sum_{i \in S} \bar{v}_i \geq (1 - 2\epsilon)g \geq (1 - O(\epsilon))OPT$.

F Details for Stochastic Knapsack with Random Profits

We first show that the $1 + \epsilon$ relaxation of the capacity constraint is necessary. Consider the following knapsack instance. The profit of each item is the same as its size. The given threshold is 1. We can see that the optimal probability is 1 if and only if there is a subset of items of total size exactly 1. Otherwise, the

optimal probability is 0. Therefore, it is NP-hard to approximate the original problem within any additive error less than 1 without violating the capacity constraint.

The corresponding exact version of the deterministic problem is to find a set of items S such that $w(S) \leq 1$ and $v(S)$ is equal to a given target value. In fact, there is no pseudopolynomial time algorithm for this problem. Since otherwise we can get an ϵ additive approximation without violating the capacity constraint, contradicting the lower bound argument. Note that a pseudopolynomial time algorithm here should run in time polynomial in the profit value(not the size). However, if the sizes can be encoded in $O(\log n)$ bits (we only have a polynomial number of different sizes), we can solve the problem in time polynomial in n and the largest profit value by standard dynamic programming.

For general sizes, we can round the size of each item down to the nearest multiple of $\frac{\epsilon}{n}$. Then, we can solve the exact version in pseudopolynomial time by dynamic programming. It is easy to show that for any subset of items, its total size is at most the total rounded size plus ϵ . Therefore, the total size of our solution is at most $1 + \epsilon$.